AMENDMENTS TO THE SPECIFICATION

Paragraphs at page 1, lines 7-18:

1. U.S. patent application, serial no. 09/693,359, entitled "Scaleable Multipath Wormhole Interconnect," Attorney Docket No. M8175US, naming John Hesse as inventor, and filed on even date herewith.

A1

2. U.S. patent application, serial number 09/693,603, entitled "Scaleable Interconnect Structure for Parallel Computing and Parallel Memory Access, Attorney Docket No. M-9051 US, naming Coke Reed and John Hesse as inventors and filed on even date herewith.

3. U.S. patent application, serial number 09/693,358, entitled "Scaleable Interconnect Structure Utilizing Quality of Service Handling, Attorney Docket No. M9051US, naming Coke Reed and John Hesse as inventors and filed on even date herewith.

4. U.S. patent application, serial number 09/693,357, entitled Scaleable Wormhole Routing Concentrator," Attorney Docket No. M-9458US, naming John Hesse and Coke Reed as inventors and filed on even date herewith.

Paragraph at page 7, lines 12-22:

A2

Referring now to FIG. 1, there is shown an interconnect structure such as was described in the Reed Patent. Three nodes are illustrated in FIG. 1. The two nodes A, 102 and B, 104 are positioned to send messages directly to a third node C,

106. Nodes B and C are on a level N of the network and node A is on a level N+1 of the network. As described in the Reed and Hesse patents, node B has priority over node A to send data to node C. When node B sends a message MB to node C on path 114, node B sends a control signal 120 informing A of the sending of MB to C so that A does not send a message MA to C in a time period that would conflict with the message MB. If there is a path from C to a target output of MA as indicated by the header of MA and there is no blocking signal from B to A then A will route MA to C on path 112. If either of these conditions does not hold, then A will send MA to a node z ~~(not shown)~~ distinct from C, with that node being on level N+1 of the network.

A3 cont.

Paragraph at page 8, lines 1-22:

Four nodes are illustrated in FIG. 2. Nodes B, C, and D are on level N of the network and node A is on level N+1 of the network. All of the output ports of the network that can be reached from node B can also be reached from nodes C and D. There are output ports that ~~than~~ can be reached from A that cannot be reached from C. For this reason, when a message travels from A to C the set of output ports that the message can reach is narrowed. Among all of the nodes in the network, node C has the highest priority to send messages to node D as node C is on the same level as node D. For this reason, when only one message M arrives at node C in a given time period, that message M can always travel to node D, and there is a path from D to a targeted output port ~~of M~~. Therefore, it is not necessary to have a buffer at node C, and by the same argument buffers are not used at any other nodes.

A3

In the Reed and Hesse patents, a message MA is not allowed to travel from A to C unless the logic associated with node A is informed that B will not send a conflicting message to C. This priority of node B over node A of sending data to Node C is enforced by a control signal from B to A. In this way, A will route MA to C provided that A "wants" to send MA to C and A is not prohibited from sending MA to C by a control signal from B to A. In case FIG. 2 is a portion of a network as described in the Reed and Hesse patents, or "Scaleable Multipath Wormhole Interconnect" patent application, node A "wants" to send MA to C provided that there is a path from C to target output port of MA as specified in the header of MA. In case FIG. 2 is a portion of the interconnect structure taught in the "Scaleable Wormhole Routing Concentrator" patent application, then node A always "wants" to send MA to C because, in the case of the concentrator, all of the outputs are acceptable output ports for MA." Alternatively the Hesse Patent took advantage of the fact that only one message could arrive at node C at a given time by allowing messages from A to C to travel to C by going through node B.

A3 cont.

Paragraph at page 10, line17 to page 11, line 2:

Logic associated with node A is capable of routing a message MA to node C. There is at least one additional node N~~, not pictured~~, so that the logic associated with node A is capable of routing MA to N. In case A routes MA to C, then logic associated with node C is capable of routing MA to nodes D and H. In this manner, the message MA can travel from A to D and the message MB can travel from B to H. The logic associated with A is incapable of routing MA to either D or

A4

A4 Cont.

H. Similarly, logic associated with B is able to route a message MB from B to C and logic associated with C can route MB to either node D or node H. So that while the message MB is able to travel from B to D or from B to H, the logic associated with node B is not capable of routing message MB to either node D or node H.

Paragraphs at page 13, lines 5-12:

3) otherwise, A sends MA to a node Z ~~(not shown)~~ distinct from C that is on the same level as A.

A5

In case two messages MA and MÁ arrive simultaneously at Node A, then one of the two messages is sent to C according to the above logic, and the remaining message is sent to [[a]] node Z, distinct from C ~~(not shown)~~. In this way, there are messages that advantageously drop down a level with the present invention that would not drop down a level in the Reed and Hesse patents. A feature of the above logic is that whenever two messages arrive simultaneously at a node, at least one of those messages will be allowed to drop to a lower level.

Paragraph at page 14, lines 1-2:

A6

2) it is guaranteed that a message sent from F to G will be sent from G to a node ~~J (not shown)~~ distinct from H.

Paragraph at page 14, lines 6-12:

Node B does not use the information contained in the bits x and y in order to determine where to send its messages; it uses information from still another control line CB(N-1) from a node on level N-1 ~~(not shown)~~ in order to determine where to send its own message. Node B uses the information in lines CEB and CFB in order to be able to send a control signal to A using the control line CBA. Node B sends a single bit z on the control line CBA. Assume that exactly one message MA arrives at node A. Then MA is sent from node A to C, provided that the bit z is zero and C lies on a path to a target of MA. The bit z is set to zero provided that either:

A7

Paragraphs at page 14, line 17 to page 15, line 3:

Node A is able to route an incoming message MA based on the header of MA and on the value of the single bit z. In case two messages MA and MÁ arrive simultaneously at A, then one of those two messages is sent to C according to the above logic, and the other message is sent to a node distinct from C such as node Z ~~(not shown)~~. A feature of the above logic is that one of the two messages MA and MÁ will be allowed to drop to C. In particular, the messages MA and MÁ are not routed to the same output port of A.

A8

It is important to note that nodes in accordance with the present embodiment are able to route messages based on one header address bit and on

A8 Cont'

control bits from lower levels. In this way the timing is the same as the timing in the Reed and Hesse patents. Importantly, with the embodiment of FIG. 7, node A is able to send a message to C in a case where node A using the logic of FIG. 6A was not able to send a message to C but instead sent its message to a node on level N+1 such as node Z.

Paragraphs at page 16, line 5 to page 17, line 3:

A9

The first condition (1) above, is discussed above, and the second condition pertains to the "cross over" case. If neither of the above conditions is satisfied, then A will send MA to a node ~~(not shown)~~ other than C such as Z, which node will be on level N+1. The case in which two messages MA and MÁ appear simultaneously at node A is handled as described above. Reading two header bits allows us to detect condition (2) above. This sometimes allows the sending from A to C of a message MA that would have stayed on the same level as A under the earlier embodiment of FIG. 6A. The reading of two header address bits requires only minor modifications to the control logic and control signals of the networks described herein and in the Reed and Hesse patents. Such modifications would be apparent to one skilled in the art of this invention and thus further description of such modifications will not be presented herein.

Note that in FIG. 7, node A can send data to node H via node C, while node F can send data to node H via node G. The control signals CFB and CBA ~~x~~

~~and z~~ enforce a priority of the transfer of data from F to H over the transfer of data from A to H.

A9 cont.

Refer now to FIG. 8. The nodes A and H of FIG. 8 are on level N-1 in column K+2. The nodes B and C at level N of column K+1 are positioned to send data directly to A and H. The nodes U and V of level N+1 in column K are able to send data directly to B, and the nodes W and X of level N+1 in column K are able to send data directly to C. The node B receives data directly from the node D at level N and sends data directly to node L at level N. The node C receives data directly from node E at level N, and sends data directly to node M at level N. ~~Not pictured in FIG. 8 is a~~ A collection R of nodes in column K ~~such that the members of R~~ are capable of sending control signals to nodes D and E. Node D uses information from a node in R ~~(not shown)~~ and node E uses the identical information from node D. The control information that node D receives from a node in R enables node D to determine if the paths from node B to node A and node H are unblocked.

Paragraph at page 20, lines 3-14:

A10

U.S. patent application, Serial No. 09/693,359, entitled "Scaleable Multipath Wormhole Interconnect," Attorney Docket No. M8175US, naming John Hesse as inventor, and filed on even date herewith, taught how to effectively use quality of service information in message headers. The teachings of U.S. Patent application, Serial No. 09/009,703, are hereby incorporated herein by reference. The techniques taught in that patent application can be effectively applied to this

A10 Cont.

invention, so that if, for example, the control signal from node D informs nodes U and V that one of node U and node V can send a message to node B, then the rules above will apply unless there is a low quality of service messages MU at node U, such that there is a path from node B to a target output port of MU and a high quality of service message MV at node V, so that at node B there is a path from node B to a target output port of MV. In this case, MV will be sent to node B and MU will be sent to a level N+1 node in column K+1. Quality of service header bits can also be used to determine the priority of messages arriving at nodes D and E.